arXiv:1404.6103v1 [cs.GT] 24 Apr 2014

An Approximate "Law of One Price" in Random Assignment Games

Avinatan Hassidim* Assaf Romm†

April 25, 2014

Abstract

Assignment games represent a tractable yet versatile model of two-sided markets with transfers. We study the likely properties of the core of randomly generated assignment games. If the joint productivities of every firm and worker are i.i.d bounded random variables, then with high probability all workers are paid roughly equal wages, and all firms make similar profits. This implies that core allocations vary significantly in balanced markets, but that there is core convergence in even slightly unbalanced markets. For the benchmark case of uniform distribution, we provide a tight bound for the workers' share of the surplus under the firm-optimal core allocation. We present simulation results suggesting that the phenomena analyzed appear even in medium-sized markets. Finally, we briefly discuss the effects of unbounded distributions and the ways in which they may affect wage dispersion.

1 Introduction

The "law of one price" asserts that homogeneous goods must sell for the same price across locations and vendors. In labor markets, it implies that workers who are equally skilled should earn the same salary. This paper makes the formal claim that even in the presence of some heterogeneity across firms and

*Bar-Ilan University, Department of Computer Science
†Harvard University, Department of Economics and HBS

workers, an approximate law remains valid.[1] While the law of one price does not sit well with empirical evidence, our model provides a useful benchmark for testing hypotheses on possible sources of wage inequality. Furthermore, the analysis elucidates how surplus related to idiosyncratic compatibility is divided among market participants.

Our basic building block is the assignment game model of Shapley and Shubik (1971). Each firm from a finite set of firms F is looking to hire exactly one worker from a finite set of workers W, in exchange for a negotiable salary. Each firm has a (possibly different) value for hiring each of the workers, and each worker has a (possibly different) reservation value for working for each of the firms. Transfers are permitted between any two parties (not just a transfer from a firm to a worker employed by it), and utility is assumed to be linear in money. Note that since transfers are freely allowed, we can describe the net productivity of each firm-worker pair by a single number.

We bolster the model by assuming that each element in the matrix of productivities is independently and identically distributed on a bounded interval. This assumption is similar in spirit to the one taken in many auction theory papers, in which bidders' valuations are assumed to be heterogeneous and determined according to some random distribution. However, unlike most of the literature on auction theory, we do not wish to study the effects of the random generation on agents' beliefs and equilibrium behavior. Instead we take a different approach and characterize the likely outcomes in a typical complete information matching market created in that manner.

The main result of this paper is that under certain regularity conditions and with high probability, in any core allocation workers are being paid approximately identical salaries, and firms receive approximately identical profits. Because of inherent heterogeneity in firms' and workers' preferences, the law of one price only holds approximately, with some workers being paid more than their peers. Nevertheless, the differences become negligible as the market becomes large, and there is no wage dispersion in a reasonably-sized random assignment game. Formally, we show that with high probability the core of an assignment game generated by i.i.d draws from a bounded random distribution with continuous and positive density is narrow, in the sense that the difference between the payoffs of agents on the same side of the market are $O\left(\frac{\log^2(n)}{n}\right)$.

[1] We choose to focus on labor markets as our main example. However, our results are also applicable to markets with heterogeneous commodities and unit demand buyers.

The approximate law of one price allows a deeper understanding of surplus division between the two sides of the market. It was shown already by Shapley and Shubik (1971) that there are core allocations which are optimal for the firms, and core allocations that are optimal for the workers. In a market with different number of firms and workers it is immediately implied that agents on the long side will get a vanishing share of the total surplus. However, in a balanced market the choice of core allocation can have compelling consequences for all agents in the market. This extends the basic economic intuition on competition in a market for homogeneous goods.

It is worthwhile to mention that the model can also be used to describe auctions of heterogeneous items with unit-demand bidders Demange *et al.* (1986). With this interpretation in mind, core allocations are equivalent to Walrasian equilibira. Our results therefore provide insight into revenue acquired by sellers under different market conditions.

From a slightly more broad point of view, this paper belongs to the theoretical literature on matching in two-sided markets. This literature gained prominence in the 1960's and early 1970's following the publication of the seminal papers by Gale and Shapley (1962) and Shapley and Shubik (1971), yet research remained mostly divided (with some notable exceptions) into two parallel strands: with and without transferable utility (i.e. money). During the 1980's, as it became clear that real-life centralized clearing houses can be immensely improved using intuitions gained in the study of marriage markets, the transferable utility strand of the literature became slightly neglected compared to its glorified non-transferable utility half-sibling. It is our belief that assignment games provide an excellent way to model decentralized markets, and that both strands of the matching theory literature can benefit from the continuous cross-fertilization.

The rest of the paper is organized as follows. Section 2 reviews the literature related to our paper. Section 3 introduces the model and the formal notation. Section 4 contains the statement and proof of the main result, and Section 5 applies the result to analyze the likely properties of balanced and unbalanced markets. Section 6 provides some simulation results and a discussion on unbounded distribution, and Section 7 concludes.

2 Related Work

Assignment games were introduced by Shapley (1955) and thoroughly analyzed by Shapley and Shubik (1971). A (non-comprehensive) list of further work on assignment games includes the study of strategic incentives Demange and Gale (1985), entry Mo (1988), convergence via decentralized processes Chen *et al.* (2012); Klaus and Payot (2012); Nax *et al.* (2013), elongation of the core Quint (1987) and its dimensions Núñez and Rafels (2008), and median stable matchings Schwarz and Yenmez (2011).

Analysis of random instances of the linear sum assignment problem was first carried by Walkup (1979), and was subsequently improved by Aldous (2001); Coppersmith and Sorkin (1999); Karp (1987). For a survey on the topic and other related literature see Krokhmal and Pardalos (2009).

While to the best of our knowledge we are the first to introduce random values into the assignment game framework, some of the consequences we present in Section 5 resemble work done on random marriage markets. Pioneered by Wilson (1972) and Knuth (1976), and extensively developed by Pittel (1989, 1992), the analysis of random preference marriage markets suggests that in a situation in which the number of men is equal to the number of women, with high probability the proposing side's (in a deferred acceptance algorithm) mean rank of partners behaves like $\log(n)$, where as the other side's mean rank of partners behaves like $\frac{n}{\log(n)}$. Recently, Ashlagi *et al.* (2013) have shown that in unbalanced random marriage markets with high probability under any stable matching the short side's mean rank of partners behaves like $\log(n)$, whereas the long side's mean rank of partners behaves like $\frac{n}{\log(n)}$. Coles *et al.* (2014) and Coles and Shorrer (2014) utilize those results to study aspects of strategic behavior in marriage markets with incomplete information. The papers most closely related to the current paper are Lee (2014) and Lee and Yariv (2014) that assume that preferences are derived from underlying cardinal utilities and study the issues of core convergence and efficiency, respectively. Finally, a different type of results on random marriage markets (and extensions thereof) are related to showing core convergence when preferences are bounded in length Ashlagi *et al.* (2011); Immorlica and Mahdian (2005); Kojima and Pathak (2009); Kojima *et al.* (2013).[2]

[2]Similar ideas were also applied by Manea (2009), Che and Kojima (2010) and Kojima and Manea (2010).

3 Model and Notation

Consider a sequence of markets $\{M^n\}_{n=1}^{\infty}$, such that each market can be described as $M^n = (F^n, W^n, \alpha^n)$, where F^n is a set of firms of size n, W^n is a set of workers of size $n + k(n)$, with $k(n) \geq 0$, and α^n being an $|F^n| \times |W^n|$ real matrix representing the value of pairs of firms and workers. We assume throughout each element of α^n is distributed i.i.d according to the cdf G which is bounded on the interval $[0, 1]$ (meaning $G(1) = 1$) and has a pdf g which is continuous and strictly positive[3].

In market M^n, the value of a coalition of firms and workers S is given by

$$\mathrm{v}(S) = \max\left[\alpha^n_{i_1 j_1} + \alpha^n_{i_2 j_2} + \cdots + \alpha^n_{i_l j_l}\right],$$

where the maximum is taken over all arrangements of $2l$ distinct agents, $i_1, \ldots, i_l \in S \cap F$, $j_1, \ldots, j_l \in S \cap W$, and $l \leq \min\{|S \cap F|, |S \cap W|\}$. An **allocation** is denoted by (μ, u, v) with μ being a matching of firms to workers and vice-versa, and u and v being payoff vectors for the firms and workers respectively. We will refer to u as firms' "profits", and to v as workers' "salaries". Formally, $\mu \colon F^n \cup W^n \to F^n \cup W^n \cup \{\emptyset\}$, and satisfies

1. $\forall f \in F^n : \mu(f) \in W^n \cup \{\emptyset\}$,
2. $\forall w \in W^n : \mu(w) \in F^n \cup \{\emptyset\}$, and
3. $\forall f \in F^n, w \in W^n : \mu(f) = w \iff \mu(w) = f$.

An allocation is a **core allocation** if no coalition can deviate and split the resulting value between its members such that each member of the coalition becomes strictly better off. We denote the set of core allocations of M^n by $C(M^n)$. Shapley and Shubik (1971) have shown that the core is a non-empty compact and convex set, and that it is elongated in the sense that there is a firm-optimal core allocation in which salaries are the lowest, and a worker-optimal core allocation in which salaries are the highest.

[3] In fact for our our results to hold we just need that the pdf is continuous near its supremum, and that it doesn't depend on n. In addition, the results also hold if there are atoms. Specifically, if there is an atom at the supremum of the distribution then whp the surplus will be n times the supremum, there will be no wage dispersion, and in a firm-optimal allocation all workers will get zero. Atoms at other places do not matter as n goes to infinity.

4 An approximate law of one price

Our main theorem states that the core is not only elongated, but also narrow, in the sense that with high probability (i.e. with probability that approaches 1) firms' profits and workers' salaries do not vary a lot.

Theorem 1. *There exists $c \in \mathbb{R}_+$ such that with probability at least $1 - O\left(\frac{1}{n}\right)$, for any $(\mu, u, v) \in C\left(M^n\right)$ we have*

1. $\forall i, j \in \{1, \dots |F^n|\} : |u_i - u_j| \leq \frac{c \log^2 n}{n}$, *and*
2. $\forall i, j \in \{1, \dots |W^n|\} : |v_i - v_j| \leq \frac{c \log^2 n}{n}$.

Let $\mathcal{G}$ be a directed bipartite graph where firms are vertices of one side, and workers are vertices on the other side. Each firm f points to the $32\frac{n+k(n)}{n}\log(n)$ workers it values the most, equivalently there is an edge (f, w) if and only if

$$|\{\alpha_{f,j} : \alpha_{f,j} \geq \alpha_{f,w}\}| \leq 32\frac{n + k(n)}{n}\log(n).$$

Note that since there are no ties in α the verbal definition coincides with the mathematical one.

Each worker that is assigned points to the firm which hired her.

Claim 2. *With probability at least $1 - 1/n$, for any set of firms $\mathcal{F}_\mathcal{S}$ with $|\mathcal{F}_\mathcal{S}| < n/10$, we have that*

$$|N(\mathcal{F}_\mathcal{S})| \geq 2\frac{n + k(n)}{n}|\mathcal{F}_\mathcal{S}|,$$

where $N(S)$ is the set of neighbors of a set of vertices S. In addition, for any set of firms $|\mathcal{F}_\mathcal{S}|$ of size $n/10$

$$|N(\mathcal{F}_\mathcal{S})| \geq 0.99(n + k(n)).$$

Proof. Consider a set of firms $\mathcal{F}_\mathcal{S}$ of size $m < n/10$, and a set of workers $\mathcal{W}_\mathcal{S}$ of size $\eta < 2\frac{n+k(n)}{n}m$. Since the edges coming out of $\mathcal{F}_\mathcal{S}$ point to independent random workers, the probability that all edges point towards workers in $\mathcal{W}_\mathcal{S}$ is at most

$$\left(\frac{\eta}{n + k(n)}\right)^{\frac{32(n+k(n))m\log n}{n}}.$$

Taking a union bound over all sets of firms of size m, and all sets of workers of size η, the probability that there is a set of firms of size m which points to $\eta < 2\frac{n+k(n)}{n}m$ workers is upper bounded by

$$\begin{aligned}
\binom{n}{m}\binom{n+k(n)}{\eta}\left(\frac{\eta}{n+k(n)}\right)^{\frac{32(n+k(n))m\log n}{n}} &\leq \\
\binom{2n}{2\frac{n+k(n)}{n}m}^2\left(\frac{2\frac{n+k(n)}{n}m}{n+k(n)}\right)^{\frac{32(n+k(n))m\log n}{n}} &= \\
\binom{2n}{2\frac{n+k(n)}{n}m}^2\left(\frac{2m}{n}\right)^{\frac{32(n+k(n))m\log n}{n}} &\leq \\
\binom{2n}{4m}^2\left(\frac{2m}{n}\right)^{32m\log(n)} &\leq \\
\left(\frac{2ne}{4m}\right)^{8m}\left(\frac{2m}{n}\right)^{32m\log(n)} &\leq 1/n^2,
\end{aligned}$$

where the second inequality uses $0 \leq k(n) \leq n$, the third uses the fact that $m < n/10$ and binomial coefficients increase until $\binom{2n}{n}$, and the fourth is an approximation of the binomial coefficient.

Using $m < n/10$, we union bound over m, and get a failure probability of at most $1/10n$.

Similarly, let $\mathcal{F}_{\mathcal{S}}$ be a set of workers with $|\mathcal{F}_{\mathcal{S}}| = n/10$. The probability that all the workers they point to lie in a set of size at most $0.99(n+k(n)$ is

$$\begin{aligned}
\binom{n}{n/10}\binom{n+k(n)}{0.99(k(n)+n)}\left(\frac{0.99(n+k(n))}{n+k(n)}\right)^{3.2n\log(n)} &= \\
\binom{n}{n/10}\binom{n+k(n)}{0.01(k(n)+n)}\left(\frac{0.99(n+k(n))}{n+k(n)}\right)^{3.2n\log(n)} &\leq \\
\left(\frac{ne}{n/10}\right)^{n/10}\left(\frac{(n+k(n))e}{(n+k(n))/100}\right)^{(n+k(n))/100}\left(\frac{0.99(n+k(n))}{n+k(n)}\right)^{3.2n\log(n)} &\leq \\
(10e)^{n/10}(100e)^{0.02n}\left(\frac{1}{100}\right)^{3n\log(n)} &\leq 1/n^2,
\end{aligned}$$

which finishes the proof. □

If $k(n) > n$, we can prove the following stronger claim, via a similar technique.

Claim 3. *With probability at least $1 - 1/n$, for any set of firms $\mathcal{F}_\mathcal{S}$, we have that*

$$|N(\mathcal{F}_\mathcal{S})| \geq 2\frac{n + k(n)}{n}|\mathcal{F}_\mathcal{S}|,$$

where $N(S)$ is the set of neighbors of a set of vertices S.

For $k(n) \leq n$, we also show a similar claim for workers, showing that for any set of workers, the set of firms which point to them is large.

Claim 4. *With probability at least $1 - 1/n$, for any set of workers $\mathcal{W}_\mathcal{S}$ with $|\mathcal{W}_\mathcal{S}| < n/10$, we have that*

$$|\{f \ : \ N(f) \cap \mathcal{W}_\mathcal{S} \neq \emptyset\}| \geq 2|\mathcal{W}_\mathcal{S}|.$$

Proof. Consider a set of workers $\mathcal{W}_\mathcal{S}$ of size $m < n/10$, and a set of firms $\mathcal{F}_\mathcal{S}$ with $|\mathcal{F}_\mathcal{S}| = \eta < 2m$. The probability that for every $f \notin \mathcal{F}_\mathcal{S}$ we have $N(f) \cap \mathcal{W}_\mathcal{S} = \emptyset$ is

$$\left(\frac{n + k(n) - m}{n + k(n)}\right)^{32(n-\eta)\log(n)}$$

Taking a union bound over sets of workers and sets of firms, the probability that there exists a set of workers of size m such that exactly η firms point to it, with $\eta < 2m$ is at most

$$\begin{aligned}
\binom{n + k(n)}{\eta}\binom{n}{n-\eta}\left(\frac{n + k(n) - m}{n + k(n)}\right)^{32(n-\eta)\log(n)} &= \\
\binom{n + k(n)}{\eta}\binom{n}{\eta}\left(1 - \frac{m}{n + k(n)}\right)^{32(n-\eta)\log(n)} &\leq \\
\binom{2n}{\eta}^2\left(1 - \frac{m}{2n}\right)^{32(n-\eta)\log(n)} &\leq \\
\binom{2n}{2m}^2\left(1 - \frac{1}{2n}\right)^{16n\log(n)} &\leq \\
\left(\frac{2ne}{2m}\right)^{4m} e^{-8m\log(n)} &\leq \\
(ne)^{4m}e^{-8m\log(n)} &\leq \\
(ne)^{4m}\left(\frac{1}{n}\right)^{8m} &\leq 1/n^3.
\end{aligned}$$

Taking a union bound over the values of m and η finishes the proof. $\square$

The following claim is the only place where we use the continuity of g, the density of the distribution G used to sample α.

Claim 5. *With probability $1 - 1/n$, if firm f points to a worker w then $\alpha_{f,w} > 1 - \frac{512\log(n)}{g(1)n}$.*

Proof. Let $\beta = 32\frac{n+k(n)}{n}\log(n)$, let τ be such that $G(\tau) = 1 - 8\beta/(n + k(n)) = 1 - \frac{256\log(n)}{n}$. The probability that out of $n + k(n)$ samples, exactly $\eta < 32\log(n)$ of them would be above τ is

$$\begin{aligned}
\binom{n+k(n)}{\eta}\left(8\beta/(n+k(n))^{\eta}\left(1-8\beta/(n+k(n))^{(n+k(n)-\eta)}\right.\right. &\leq \\
\binom{n+k(n)}{\beta}\left(8\beta/(n+k(n))^{\beta}\left(1-8\beta/(n+k(n))^{\frac{n+k(n)}{2}}\right.\right. &\leq \\
\left(\frac{e(n+k(n))}{\beta}\right)^{\beta}\left(8\beta/(n+k(n))^{\beta}\left(1-8\beta/(n+k(n))^{\frac{n+k(n)}{2}}\right.\right. &= \\
(8e)^{\beta}\left(1-8\beta/(n+k(n))^{\frac{n+k(n)}{2}}\right. &\leq \\
(8e)^{\beta}e^{-4\beta} &= \\
e^{-(4-\log 8)\beta} < 1/n^3. &
\end{aligned}$$

Taking a union bound over all possible values of $\eta < \beta$ and n firms gives the result.

Since g is continuous, there is an interval $[a, 1]$ in which $g(x) \geq g(1)/2$. We choose $\tau \geq 1 - \frac{512\log(n)}{g(1)n}$. $\square$

We now use the graph and its properties to lower bound the profits of the firms.

Claim 6. *If a firm i points to worker that's unemployed then $u_i \geq 1 - \frac{512\log(n)}{g(1)n}$. Moreover, if firm i points to a worker j, and j points to i' then $u_i \geq u'_i - \frac{512\log(n)}{g(1)n}$.*

Proof. If i points to j which is unemployed, i could hire j and pay her nothing. Therefore, it must be that $u_i \geq \alpha_{i,j} \geq 1 - \frac{512\log(n)}{g(1)n}$.

For the other case, i points to j, and j points to i'. since $\alpha^n_{i',j} \leq 1$, it must be that $v_j \leq 1 - u_{i'}$. But since it is an equilibrium

$$u_i \geq \alpha^n_{i,j} - v_j \geq \alpha^n_{i,j} - 1 + u_{i'} \geq 1 - \frac{512\log(n)}{g(1)n} - 1 + u_{i'} = u_{i'} - \frac{512\log(n)}{g(1)n},$$

where we used Claim 5 □

Note that in particular if there is a path of length k from a firm i to a firm i' then $u_i \geq u'_i - \cdot k$, and if there is a path of length k from firm i to an unemployed worker than $u_i \geq 1 - \frac{512\log(n)}{g(1)n}k$.

Proof of Theorem 1. Consider any core outcome and two firms i, i'. We want to show that $u_i \geq u_{i'} - \frac{1024\log^2(n)}{g(1)n}$.

Let $S_1 = N(i)$ be the set of workers which i points to. If there is $w \in S_1$ such that $\mu(w) = \emptyset$, then by Claim 6 $u_i \geq \alpha_{i,w} \geq 1 - \frac{512\log(n)}{g(1)n}$, and since $u_j \leq 1$, we are done.

If for every $w \in S_1$ we have $\mu(w) \neq \emptyset$, then $|N(S_1)| = |S_1|$. In this case, let $S_2 = N\left(N(S_1)\right)$, and by Claim 2 (or Claim 3 if $k(n) > n$) we have $|S_2| \geq 2 \cdot |S_1|$. If there is $w \in S_2$ such that $\mu(w) = \emptyset$, there is a path of length 2 to an unemployed worker, and $u_i \geq 1 - \frac{1024\log(n)}{g(1)n}$.

Continuing this by induction, if for some $w \in S_k$ we have $\mu(w) = \emptyset$ then $u_i \geq 1 - k\frac{512\log(n)}{g(1)n}$. Otherwise, we define $S_{k+1} = N\left(N(S_k)\right)$, and get that $|S_{k+1}| \geq 2|S_k|$. If $k(n) > n$, We continue with this process until we get to an unemployed worker and we are done in $\log(n)$ steps.

If $k(n) \leq n$, we can continue for at most $log(n) - 3$ steps, after which we would be left with a set of firms $\mathcal{F}_\mathcal{S}$ of size $n/8$.

We now start with the target firm i'. There is a worker w that points to it. According to Claim 4, there are at least two firms that point to w. Now there are two workers who point to those firms, four firms that point to them, etc. We continue this for $\log(n) - 4$ steps until we get to a set of workers $\mathcal{W}_\mathcal{S}$ of size at least $n/16$. However, invoking the second part of Claim 2, the set $\mathcal{F}_\mathcal{S}$ points to at least $0.99(n + k(n))$ workers, and since $|\mathcal{W}_\mathcal{S}| > 0.01(n + k(n))$ by the pigeonhole principle there is an edge from $\mathcal{F}_\mathcal{S}$ to $\mathcal{W}_\mathcal{S}$, and a path of length at most $2\log(n) - 6$ from i to i', which implies the theorem for firms.

We now show the second part of the theorem, namely that workers get similar salaries. If the market is balanced ($k(n) = 0$), we can apply the same argument. Else, let w be a worker, and we show that $v_w \leq \frac{1024\log^2(n)}{g(1)n}$. If w

is unemployed then $v_w = 0$ and we are done. Else, let f be the firm which hires w.

If $k(n) > n$, the firm f has a path of length $\log(n)$ to an unemployed worker and therefore $u_f \geq 1 - \frac{512\log^2(n)}{g(1)n}$ and therefore $v_w \leq \frac{512\log^2 n}{n}$.

If $k(n) \leq n$, let w' be an unemployed worker, which has a firm f' pointing to it. Since for $k(n) \geq n$, any two firms have a path of length at most $2\log(n) - 6$ between them we have

$$u_f \geq u_{f'} - (2\log(n) - 6)\frac{512\log(n)}{g(1)n},$$

and since f' points to an unemployed worker $u_{f'} \geq 1 - \frac{512\log(n)}{g(1)n}$. Therefore $u_f \geq 1 - (2\log(n) - 4)\frac{512\log(n)}{g(1)n}$. But since $u_f + v_w \leq 1$, it must be that $v_w \leq \frac{1024\log^2(n)}{g(1)n}$ as required. □

5 Applications

5.1 Balanced markets

We first deal with the case of $k(n) \equiv 0$. In this case we will show that whp the firm-optimal core allocation gives almost all the surplus to the firms. From symmetry this will imply that the core is "long" in the sense that different market mechanisms can lead to very different divisions of the surplus.

Corollary 7. *In a balanced market, with probability* $1 - O\left(\frac{1}{n}\right)$*, under the firm-optimal core allocation the workers' share of the surplus is* $O\left(\frac{\log^2(n)}{n}\right)$.

Proof. Under the firm-optimal core allocation at least one worker must get exactly 0 (or else we can substract some small $\epsilon > 0$ from each worker's salary and remain with a new core allocation which is even better for the firms). Theorem 1 then immediately implies that the sum of workers' salaries is $O(\frac{\log^2(n)}{n})$, whereas routine arguments from the theory of random assignment problems prove that the total surplus is $n - o(n)$. □

For the benchmark case of uniform distribution we can also provide tight bounds for the sum of workers' salaries under the firm-optimal core alloca-

tion.[4]

Theorem 8. *In a balanced market with uniform distributions, $G = U[0,1]$, let $\psi^F(M^n) = \left(\mu^{F,n}, u^{F,n}, v^{F,n}\right)$ be the firm-optimal core allocation. Then*

1. $E\left[\sum_{i\in\{1,\ldots,n\}} v_i^{F,n}\right] \leq \log(n)$, *and*
2. $E\left[\sum_{i\in\{1,\ldots,n\}} v_i^{F,n}\right] = \Omega\left(\log(n)\right)$.

Proof. For the upper bound, consider a variant of the approximation algorithm suggested by Crawford and Knoer (1981) to solve a generalized version of the assignment game, in which firms are ordered from 1 to n, and at each round only the lowest-number firm that still wants to propose actually proposes. What is the expected number of proposals between the first proposal of firm k and the first proposal of firm $k+1$? Assume that it is now firm f's turn to propose, and its previous aspiration level (i.e. the maximal utility it would get by giving some worker her current salary) was u_i. If for some unmatched $w \in W^n$ we have $\alpha_{f,w}^n \in [u-\epsilon, u)$ then firm f will propose to that worker within at most k rounds (or otherwise it will temporarily match to some other worker and be replaced by another proposing firm). Therefore the probability not proposing to some unmatched worker is above below by

$$1-(1-\epsilon)^{n-k+1} \approx (n-k+1)\epsilon,$$

and so the expected number of steps until the first proposal of firm $k+1$ is at most $\frac{1}{(n-k+1)\epsilon}$, and the the expected raise in workers' salaries is $\frac{1}{n-k+1}$. Summing over k we get that the expected sum of workers' salaries is at most

$$1+\frac{1}{2}+\cdots+\frac{1}{n-1}+\frac{1}{n} \approx \log(n).$$

For the lower bound a similar exercise can be taken, but taking into account that whp all workers' temporary salaries never exceed $1-\frac{c\log^2(n)}{n}$. The details of the proof are omitted. □

[4] In fact, the same proof strategy works, with minor modifications, for any distribution with continuous and strictly positive density. Furthermore, it also applies to unbalanced markets in which the number of firms is smaller than the number of workers.

5.2 Unbalanced markets

The same reasoning that works for balanced markets can be applied to unbalanced markets. Our next result states that even slightly unbalanced markets are likely to give all the surplus to the short side.

Corollary 9. *In an unbalanced market, with probability* $1 - O\left(\frac{1}{n}\right)$*, under any core allocation all agents on the long side get* $O\left(\frac{\log^2(n)}{n}\right)$*, and the long side's share of the surplus is* $O\left(\frac{\log^2(n)}{n}\right)$ *as well.*

Proof. Immediate from Theorem 1 and the fact that under any core allocation at least one agent on the long side remains unmatched. □

6 Simulations

While the bound established in Theorem 1 converges to zero quite rapidly, it is still worth while to see how powerful is the phenomenon described. Indeed, as we stress in the concluding section, we expect the wage dispersion to behave like $O\left(\frac{\log(n)}{n}\right)$. In this section we present results of computerized simulations that demonstrate how quick is the contraction of payoffs' dispersion, and how this affects the market. We then use more simulations to explore payoff dispersion in the presence of unbounded distributions. Each figure is based on averaging 400 trials for each data point.

We first focus on the benchmark case of uniform $[0, 1]$ distribution, and study wage dispersion in balanced markets under the firm-optimal core allocation. The right panel of Figure 1 shows that in this case both the mean salary of workers and the maximum salary any worker gets (which is also the maximal difference between any two workers' salaries) approach zero very quickly. The left panel of the same figure exemplifies the fact that the core in balanced markets is long, as suggested by Theorem 7.

In unbalanced markets we expect the core to be much more narrow, per Theorem 9. The left panel of Figure 2 shows that when the number of workers is only slightly larger than the number of firms, both the workers' mean salary and the maximum salary any worker gets, approach zero rapidly, even under the worker-optimal core allocation. Furthermore, as the right panel demonstrates, in this case the workers' share in the surplus approaches 0, even under the worker-optimal core allocation. Finally, Figure 3 parallels

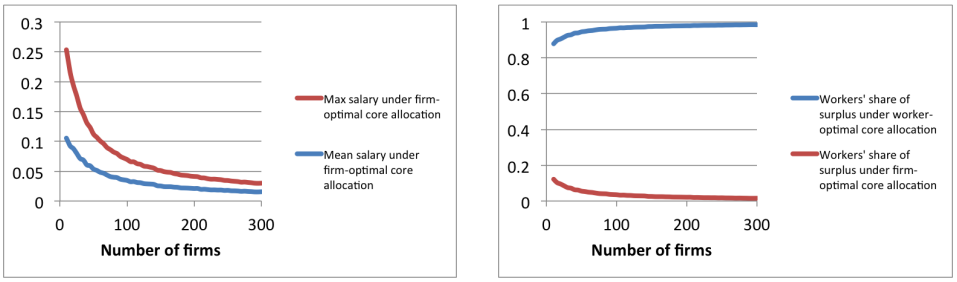

Figure 1: Balanced markets (n firms, n workers), distribution: $U[0,1]$

Figure 4 of Ashlagi *et al.* (2013), and depicts workers' share of the surplus when the number of workers is constant at 50, and the number of firms varies from 20 to 80.

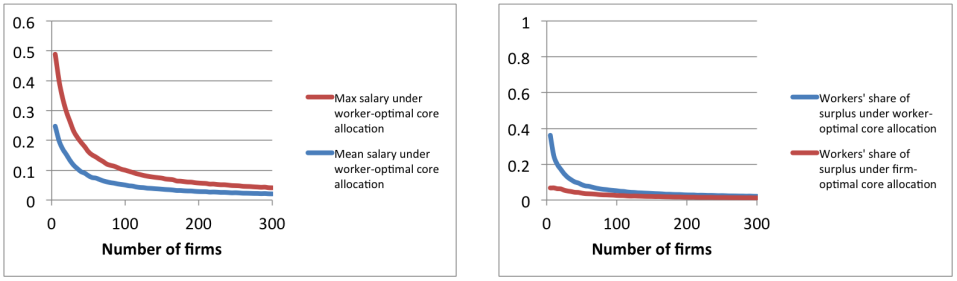

Figure 2: Unbalanced markets (n firms, $n+1$ workers), distribution: $U[0,1]$

6.1 Unbounded distributions

A very interesting question is what happens to our results when the distribution G is unbounded. A rough intuition for this case is that unbounded distributions with a heavy tail may create two types of outliers in the market. "Good" outliers are agents that are highly productive compared to others, such that agents from the other side fiercely compete on being matched with them. These agents share a significant portion of the surplus they help create, and if they are common enough, they may offset other forces that would otherwise squeeze the surplus from their side (such as an adversarial core allocation, or slight imbalance in favor of the other side of the market). "Bad" outliers are agents that are so unproductive in comparison with their peers, that being matched with them becomes the effective reservation value for

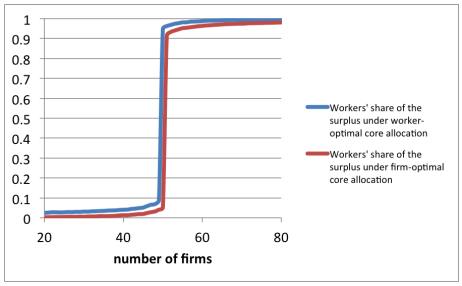

Figure 3: Unbalanced markets (50 workers), distribution: $U[0,1]$

agents on the other side. If for any reason there are many such agents on one of the sides, then the residual market behaves as if it was unbalanced, and even mildly productive agents that are on the same side as the less-productive agents earn a sizable part of their contribution to the total surplus.

Figure 4 describes workers' salaries in a balanced market governed by two distributions: Exponential and Weibull with parameter 0.25 (a heavy-tailed distribution). As expected, we can see that the maximal salary does not go to 0, and behaves more like $\log(n)$. In Figure 5 we can see the workers' share in the total surplus for the same markets.

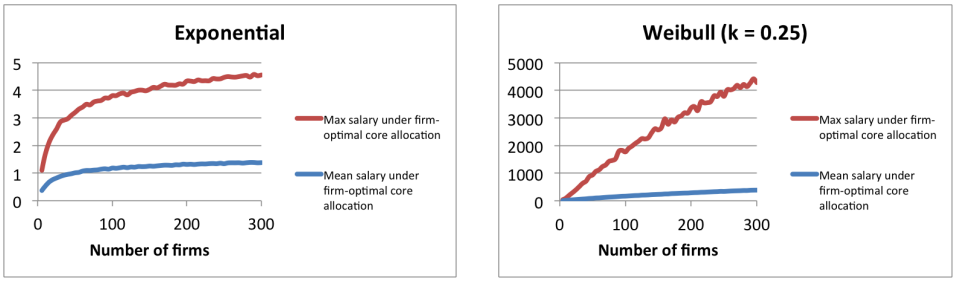

Figure 4: Balanced markets, workers' salaries

We conclude this section by noting that it is relatively easy to establish analytically that Theorem 1 does not hold for the case of an exponential distribution:

Proposition 10. *If G is the exponential distribution and the market is balanced and governed by the firm-optimal core allocation, then whp there are two workers i and j such that $|v_i - v_j| = \Omega(\log(n))$.*

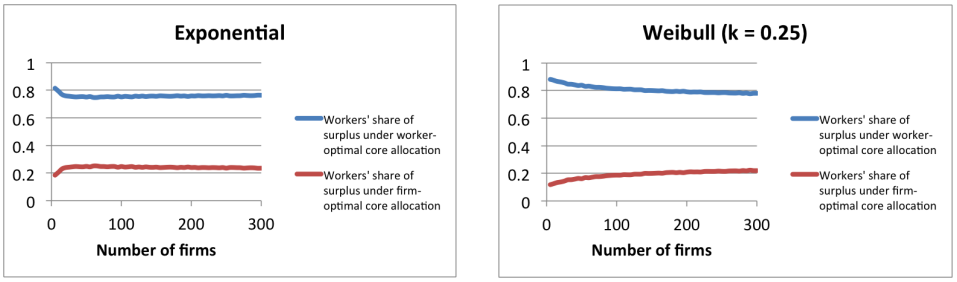

Figure 5: Balanced markets, workers' share of surplus

Proof. Let p^n denote the probability that for a given firm $f \in F^n$ there exists a worker $w \in W^n$ such that $\alpha^n_{f,w} > 1.1\log(n)$ and $\max_{w' \in W^n \setminus \{w\}} \alpha^n_{f,w'} < \log(n)$. Then

$$p = n \cdot e^{-1.1\log(n)} \cdot \left(1 - e^{-\log(n)}\right)^{n-1} = \frac{1}{n^{0.1}} \cdot \left(1 - \frac{1}{n}\right)^{n-1} \approx \frac{1}{en^{0.1}}.$$

This specifically implies that whp there are $\Omega\left(n^{0.8-\epsilon}\right)$ firms that meet the above condition. If any of those two firms happen to have the same worker being the outlier, then this worker must get paid at least $0.1\log(n)$ under any core allocation. Since there are $\Omega\left(n^{1.6-2\epsilon}\right)$ pairs, then we get that there are many workers that get paid $\Theta\left(\log(n)\right)$. Finally, at least one worker's salary is 0 under the firm-optimal core allocation, and so we are done. □

6.2 The role of outliers

To give some intuition to the way wage dispersion looks, and to the role of "Good" and "Bad" outliers, Figure 6 presents the wage distribution of workers in a single instance of a firm-optimal balanced markets ($n = 1000$). In the left pane, G is the uniform distribution, while in the right pane G is the exponential distribution.

It is not surprising that salaries are higher in the exponential distribution, and that the total surplus is higher. Two other conclusions one can draw are:

1. The wage distribution in the uniform distribution is balanced, where in the exponential distribution it has a tail to the right

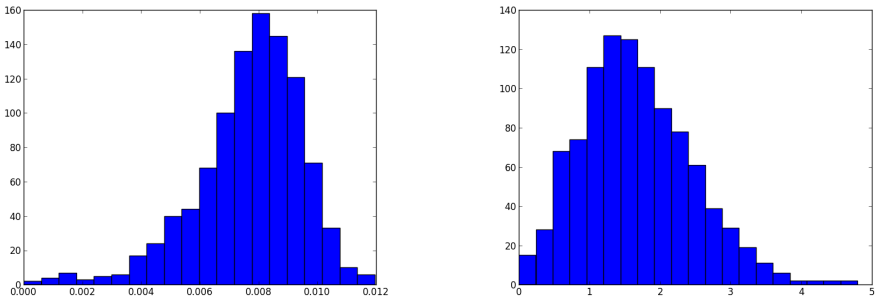

Figure 6: Balanced markets, workers' salaries

2. Although the wages in the exponential distribution have a tail, it is not heavy, and most of the salary comes from the majority of the population (there is no 1% which takes the economic pie). We see this as evidence that the contribution to the workers' salaries come from the existence of "Bad" outliers and their effect on the population.

In Figure 6 it is hard to see the effect that "Bad" outliers have in the uniform distribution. To better understand the effect, Figure 7 depicts the relationship between the sum of salaries of the workers in a firm-optimal outcome, and the quality of the "Worst" workers. The left pane shows the the dependency of the sum of salaries on the quality of the "Worst" worker (minimum over workers of the maximum value that worker has with a firm), and is based on 100 trials with $n = 100$. In the right pane, we present the effect of the second "Worst" worker. To this end, we sampled 100 random markets, for which the quality of the worst worker was 0.95 ± 0.00001, and depicted the dependency of the sum of salaries in the second "Worst" worker.[5]

Two observations about the graphs are

1. The dependency between the total salary and the quality of the worst worker becomes linear (with a slope of n) as the quality of the worker becomes lower. The reason is that it serves as a "reservation value" for the market, and reducing its quality by ϵ increases (almost) all wages by (almost) ϵ. We note that when the quality of the worst worker is

[5]We did rejection sampling, that is we sampled random markets, and kept them if the quality of the worst worker was between 0.94999 and 0.95001. This required us around a million samples.

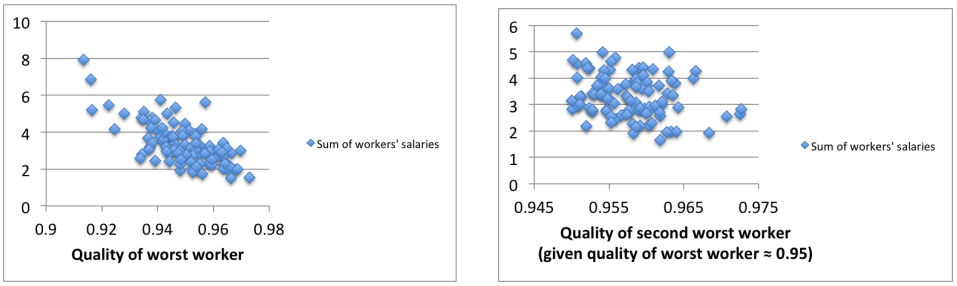

Figure 7: Balanced markets (n firms, n workers), distribution: $U[0,1]$

high, sometimes other workers serve as reservation value, and therefore the slope is not linear.

2. The second "Worst" worker also has a linear effect on the market, although here it is not as prominent.

We note that fitting the first graph to a linear approximation gives an R^2 value of 0.48, and an R^2 value of 0.07 for the second graph. The second highest value of the "Worst" worker also has an effect on the market. We conjecture that the sum of salaries that the workers get is greatly affected by the $\log(n)$ "Worst" workers.

We also note that the quality of the "Worst" firm (minimum over the firms of the maximal value this firm has with a worker) has little effect on the sum of salaries the workers get.

7 Conclusion

As hinted in the introduction, there is an abundance of real-world data suggesting that not only does wage inequality exist, but that it persists across nations, industries and even within firms (see, e.g. Mortensen (2005) and references therein). While investment and variation in human capital explain a significant portion of the differences in income, there are residual discrepancies that are not so easily explained. Economists have considered many possible approaches to this issue (for example, compensating differentials, motivating effort, complementarities, search frictions, and more) and usually carried an analysis under extremely restrictive assumptions on match value heterogeneity. Our contribution provides tools that may allow interacting those explanations within a more realistic framework.

To give but two examples, our model inherently allows for compensating differentials (by changing reservations values and then applying the analysis), and it is relatively straightforward to add human capital or quality of firms, as long as the production function is separable in those two factors. If, say, workers' quality is captured by a linear (common) contribution to each element in the relevant column, it is easy to show that each worker will get its contribution above the lowest type, and the rest of the surplus will be distributed according to similar principles to those described in this paper. Non-separable production function (that is, complementarities) require more work because the proof must take into account that assortative matching will take place to some extent (see also below).

There are several open questions that present future challenges for our model. The first is providing a tighter bound on wage dispersion. The proof of Theorem 1 used the upper $\frac{\log(n)}{n}$ fraction in order to create an expander, and used it to find paths between different firms. One could try to do so based on the upper $\frac{c}{n}$ fraction, and try to get the correct bound of $\frac{\log(n)}{n}$. There are a couple of obstacles to this approach. The first, is that for some firms using the top $\frac{c}{n}$ edges would already create a loss of $\frac{\log(n)}{n}$ (see Claim 5). A potentially bigger obstacle is that the edges from workers to firms are not random, and are chosen by the allocation. Our proof treats them as adversarial, but this requires us to use more edges from firms to workers (see Claim 4).

A different direction is studying the induced salary distribution for the exponential distribution and for general unbounded distributions. As mentioned in Section 6, unbounded distributions call for new intuitions that may offset the current analysis, and studying them is interesting. Specifically for the exponential distribution, we believe that the sum of workers' salaries in a balanced market governed by the firm-optimal core allocation is $\Theta\left(n \log\left(\log(n)\right)\right)$. The (rough) intuition is that the min max over the productivity matrix' entries behaves like $\log(n) - \log\left(\log(n)\right)$, whereas solving the assignment problem would give a mean surplus of $\log(n)$. Moreover, in expectation only $\frac{n}{\log(n)}$ of the agents would have a value above $\log(n) + \log(\log(n))$, and since the maximal value in the matrix is $2\log(n)$ their total salary can not exceed $2n$.[6]

Finally, it is also interesting to analyze correlated random matrices. As

[6]The procedure used in the proof of Theorem 10 only used the "good" outliers, and we think that most of the salary comes due to the existence of bad outliers - see Figure 6.

mentioned above, adding correlation can significantly distort some of our results. Adding separable noise with agent-specific mean can easily be dealt with, however interaction variables (such as Cobb-Douglas production function) are more tricky. With sufficient agents, we can get near-efficiency on both dimensions (the common element and the pair-specific productivity), but determining who gets what is more involved.

References

Aldous, David J. 2001. The ζ (2) limit in the random assignment problem. *Random Structures & Algorithms*, **18**(4), 381–418.

Ashlagi, Itai, Braverman, Mark, and Hassidim, Avinatan. 2011. Matching with couples revisited. *pp. 335–336 in Proceedings of the 12th ACM conference on Electronic commerce*. ACM.

Ashlagi, Itai, Kanoria, Yashodhan, and Leshno, Jacob D. 2013. Unbalanced random matching markets. *pp. 27–28 in ACM Conference on Electronic Commerce*.

Che, Yeon-Koo, and Kojima, Fuhito. 2010. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, **78**(5), 1625–1672.

Chen, Bo, Fujishige, Satoru, and Yang, Zaifu. 2012. *Decentralized market processes to stable job matchings with competitive salaries*. mimeo.

Coles, Peter, and Shorrer, Ran. 2014. Optimal truncation in matching markets. *Games and Economic Behavior*. forthcoming.

Coles, Peter, Gonczarowski, Yannai, and Shorrer, Ran. 2014. *Strategic Behavior in Unbalanced Matching Markets*. mimeo.

Coppersmith, Don, and Sorkin, Gregory B. 1999. Constructive bounds and exact expectations for the random assignment problem. *Random Structures & Algorithms*, **15**(2), 113–144.

Crawford, Vincent P, and Knoer, Elsie Marie. 1981. Job matching with heterogeneous firms and workers. *Econometrica: Journal of the Econometric Society*, 437–450.

Demange, Gabrielle, and Gale, David. 1985. The Strategy Structure of Two-Sided Matching Markets. *Econometrica*, **53**(4), 873–888.

Demange, Gabrielle, Gale, David, and Sotomayor, Marilda. 1986. Multi-item auctions. *The Journal of Political Economy*, 863–872.

Gale, David, and Shapley, Lloyd S. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, **69**(1), 9–15.

Immorlica, Nicole, and Mahdian, Mohammad. 2005. Marriage, honesty, and stability. *pp. 53–62 in Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics.

Karp, Richard M. 1987. An upper bound on the expected cost of an optimal assignment. *pp. 1–4 in Discrete algorithms and complexity: Proceedings of the Japan-US Joint Seminar*, vol. 15. Academic Press, New York.

Klaus, Bettina, and Payot, Frédéric. 2012. *Paths to stability in the assignment problem*. mimeo.

Knuth, Donald Ervin. 1976. *Mariages stables et leurs relations avec d&autres problèmes combinatoires*. Presses de l'Université de Montréal.

Kojima, Fuhito, and Manea, Mihai. 2010. Incentives in the probabilistic serial mechanism. *Journal of Economic Theory*, **145**(1), 106–123.

Kojima, Fuhito, and Pathak, Parag A. 2009. Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review*, **99**(3), 608–27.

Kojima, Fuhito, Pathak, Parag A., and Roth, Alvin E. 2013. Matching with Couples: Stability and Incentives in Large Markets. *The Quarterly Journal of Economics*, **128**(4), 1585–1632.

Krokhmal, Pavlo A, and Pardalos, Panos M. 2009. Random assignment problems. *European Journal of Operational Research*, **194**(1), 1–17.

Lee, SangMok. 2014. *Incentive compatibility of large centralized matching markets*. mimeo.

Lee, SangMok, and Yariv, Leeat. 2014. *On the Efficiency of Stable Matchings in Large Markets*. mimeo.

Manea, Mihai. 2009. Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics*, **4**(2), 165–197.

Mo, JiePing. 1988. Entry and Structures of Interest Groups in Assignment Games. *Journal of Economic Theory*, **46**(1), 66 – 96.

Mortensen, Dale. 2005. *Wage dispersion: why are similar workers paid differently*. The MIT Press.

Nax, Heinrich H, Pradelski, Bary SR, and Young, H Peyton. 2013. *Decentralized dynamics to optimal and stable states in the assignment game*. mimeo.

Núñez, Marina, and Rafels, Carles. 2008. On the dimension of the core of the assignment game. *Games and Economic Behavior*, **64**(1), 290–302.

Pittel, Boris. 1989. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, **2**(4), 530–549.

Pittel, Boris. 1992. On likely solutions of a stable marriage problem. *The Annals of Applied Probability*, 358–401.

Quint, Thomas. 1987. *Elongation of the Core in an Assignment Game*. Tech. rept. DTIC Document.

Schwarz, Michael, and Yenmez, M Bumin. 2011. Median stable matching for markets with wages. *Journal of Economic Theory*, **146**(2), 619–637.

Shapley, Lloyd S. 1955. *Markets as cooperative games*. Rand Corporation.

Shapley, Lloyd S, and Shubik, Martin. 1971. The assignment game I: The core. *International Journal of Game Theory*, **1**(1), 111–130.

Walkup, David W. 1979. On the expected value of a random assignment problem. *SIAM Journal on Computing*, **8**(3), 440–442.

Wilson, LB. 1972. An analysis of the stable marriage assignment algorithm. *BIT Numerical Mathematics*, **12**(4), 569–575.